REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE		3. DATES COVERED *(From - To)*
15-08-2011	Final Report		01-03-2008 to 31-05-2011

4. TITLE AND SUBTITLE
Geometric Methods for ATR: Shape Analysis, Object/Image
Metrics, Shape Reconstruction, and Shape Statistics

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-08-1-0118

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Dr. Peter F. Stiller

Department of Mathematics, MS 3368
Texas A&M University
College Station, Texas 77843-3368

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Texas A&M Research Foundation
Valley Park Center
400 Harvey Mitchell Parkway S., Suite 100
TAMU 3578
College Station, Texas 77845-3578

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 North Randolph Street
Arlington, Virginia 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-OSR-VA-TR-2012-0233

12. DISTRIBUTION / AVAILABILITY STATEMENT

DIST. A

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This report summarizes the work done over the approximately three years of this project. Our general focus has been on the employment of advanced mathematical techniques from algebraic and differential geometry to solve various problems associated with object recognition. The principal framework has been the theory of object/image equations and object/image metrics which we have developed and exploited in multiple contexts and for multiple sensor types. This included not only issues of recognition and target identification, but problems related to 3D reconstruction, recovery of shape from motion, feature extraction, and statistical shape analysis.

15. SUBJECT TERMS
object/target recognition, object/image equations and metrics, 3D shape reconstruction, affine transformations, shape theory, shape coordinates, invariants, Grassmannians, GPCA

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT U	**b. ABSTRACT** U	**c. THIS PAGE** U	UU	13	Dr. Peter F. Stiller
					19b. TELEPHONE NUMBER *(include area code)* 979-862-2905

Final Report
Air Force Office of Scientific Research
Grant No. FA9550-08-1-0118

"Geometric Methods for ATR: Shape Analysis, Object/Image Metrics, Shape Reconstruction, and Shape Statistics"

Dr. Peter F. Stiller, Principal Investigator
Professor of Mathematics and Computer Science
Department of Mathematics, MS 3368
Texas A&M University
College Station, Texas 77843-3368
stiller@math.tamu.edu   (979) 862-2905

Ms. Amanda Johnson, Project Administrator II
Texas A&M Research Foundation
Valley Park Center
400 Harvey Mitchell Parkway S., Suite 100
TAMU 3578
College Station, Texas 77845
ajohnson@rf-mail.tamu.edu   (979) 458-4730

August 15, 2011

Objectives: This three-year research effort was conducted at Texas A&M University by the principal investigator, Dr. Peter F. Stiller, and a number of graduate research assistants.

We begin this report by reviewing the project's objectives as outlined in the original proposal's Statement of Objectives:

The principal investigator, together with several graduate student assistants, in collaboration with scientists and engineers at the Air Force Research Laboratory (AFRL/RYAT), will conduct research on invariant shape theoretic methods for automatic target recognition. This research will include problems covering a variety of sensor models and modalities. Our approach will make use of relatively small numbers of sensed features associated with locations on the target geometry, although we plan to consider larger point clouds, or even continuous geometry in some cases. We expect to achieve target identification across all poses and, if desired, at varying scales without resorting to exhaustive template matching. To do this, we plan to make use of the emerging mathematical theory of shape to characterize internal relationships among features, independent of relevant transformations like rotation, translation, and scale. These characterizations should also turn out to be independent of the coordinate systems used to record the target or image feature locations. We plan to derive necessary and sufficient conditions for the shape of a 3D configuration of features to match a given 2D shape consisting of similar features. These matching equations (object/image equations) will be invariant to transformations of the 3D features (rotation, translation, scale etc.), transformations of the 2D features, and changes in the coordinate systems used to express the features in 3D or 2D. Moreover, they will be independent of the sensor parameters. Because of their universal invariance, these equations should provide the most robust matching criteria at least as far as object and image feature geometry is concerned. In addition, we expect to use this approach to create natural metrics that express the similarity or dissimilarity of two objects, two images, or an object-image pair. Zero value for these metrics will mean matching up to the relevant transformations and/or projections.

Once we understand the contribution of a single image toward the recognition or recovery of the geometry/shape of the object for different sensors, it will be easier to develop methods to integrate information from multiple images taken by uncalibrated, distributed sensors of varying types, or to make use of a series of images taken by a single sensor of a moving object. It will also be easier to understand and create flexible algorithms adapted to situations where the objects are not rigid, but deformable, as is the case with many of the recognition problems in biometric or medical applications (e.g. face recognition, detecting heart or tissue anomalies, gait recognition, etc.) We plan to investigate a range of problems related to all these issues.

Our ultimate goal in all cases is to exploit shape analysis to improve on and develop new algorithms for object/target recognition. After analyzing the shape spaces for various sensor types and features sets, and understanding the relations between them, we will focus on a number of important problems. This includes three major topics: 1) using the natural metrics on the shape spaces (which provide a distance (difference) measure between two object configurations or two image configurations) to find natural the object-image metrics that express the distance (failure to match) between an object and image pair. These metrics will be pose and view invariant and will be expressed in coordinate free terms. Moreover, such metrics will provide a basis for efficient hashing schemes to do target identification quickly and will support a rigorous approach to error analysis in ATR; 2) exploring reconstruction of an object's 3D shape from

2D or 1D sensed information, either from multiple sensors or from multiple images of a moving object; 3) investigating statistical issues surrounding random shapes, distributions of shapes, deformations of shapes, and noise in object recognition and ATR systems. Far ranging applications of this work could include biometric, medical, navigational, and mapping (robotic reconstruction of an environment) recognition or reconstruction problems.

A summary list of the research topics that are included in this proposal appears below:

Proposed Tasks and Problems

1. Shape Spaces and Object/Image Relations
2. Extending the O/I Formulation to Other Feature Sets and Other Sensor Models
3. Recognizing Articulated Objects
4. Metrics
5. Geometric Hashing
6. Unlabeled Matching – Permutation and Correspondence Problem
7. Shapelets
8. Point Clouds
9. 3D Shape Reconstruction and Shape from Motion
10. Probability and Statistics of Shape
11. Noise Analysis and Performance Prediction
12. Technology Transfer
13. Other Applications

Status of Effort: (Period of Performance 01/03/08 to 31/05/11.)

At the time of this writing the research effort supported by this grant has been completed. This report covers the full period of performance.

Introduction

Recall that in previous AFOSR sponsored work we were able to achieve several important results, including the understanding, development, and analysis of a global approach to invariants and object/image equations in the generalized weak perspective (affine) case. That work also included our initial construction of a new class of discrimination metrics that are generalizations of the classical Procrustes metric of statistical shape theory. In the first instance, we provided a complete dictionary between the old algebraic approach to invariants and the new, more geometric, global approach. This was worked out completely in the generalized weak perspective case and appears in our paper "Object/Image Relations, Shape Spaces, and Metrics" and in a book chapter entitled "Object-Image Metrics for Generalized Weak Perspective Projection," in *Statistics and Analysis of Shapes,* edited by Hamid Krim and Anthony Yezzi, Jr. and published by Birkhauser. This new approach creates a geometric framework for discrimination theory and a more robust approach to recognition. Some of the main ideas and their application to the full perspective (optical) case were presented in our paper, "Global Invariant Methods for Object Recognition" described in a previous report. Additional results on this topic appear in our paper "Recognizing point configurations in full perspective," which was written jointly with our graduate student Kevin Abbott for the Electronic Imaging Conference, Vision Geometry XV.

Overall our global approach provides a way to explore the behavior of recognition algorithms when dealing with multiplicities or geometric degeneracies (which cannot be handled with other methods). The difficulty in using the classical numerical invariants for this purpose is that they are only rational functions on the appropriate quotient variety. As such,

they are not always defined. This leads to serious numerical problems in any algorithm based on these invariants. To remedy these problems, we succeeded in replacing these invariants by points in a Grassmann manifold in the weak perspective case, or by certain geometric objects, namely toric sub-varieties of Grassmannians, in the full perspective case. The object/image equations become the expression of certain incidence relations in the weak perspective case or, in the full perspective case, certain "resultant-like" expressions for the existence of a non-trivial intersection of the toric sub-varieties with certain Schubert varieties in the Grassmannian. This "global" approach to invariants is providing more robust object recognition algorithms. Moreover, by representing the relevant shape spaces as varieties embedded in projective space, we can endow each shape space with a metric by restricting the standard Fubini-Study metric. These ideas are discussed in our paper "Object Recognition from a Global Geometric Perspective - Invariants and Metrics." This approach produces a natural metric on both the object and the image space that can be exploited to create an effective discrimination theory (i.e. a meaningful notion of "distance" between objects, between images, and between an object and an image.) Finally, several new directions have emerged from our work. These directions have been incorporated into our research and include the study of object/image equations for unordered point features to facilitate point cloud matching, research on object/image equations with parameters to handle articulation of objects, the investigation of invariant point to surface matching, 3D reconstruction from motion, and the statistics of shape for noise analysis. Progress on these will expand the recognition power of our approach and its applicability to Air Force problems.

Accomplishments/Findings:

We report below (in summary form) on several significant areas of progress. Details can be found in the listed papers.

1. 3D shape reconstruction from motion.

(Details of our work in this area can be found in our paper, "Aspects of 3D Reconstruction," which is cited in the publication list below.)

The ability to reconstruct the three dimensional (3D) shape of an object from multiple images of that object is an important step in certain computer vision and object recognition tasks. The images in question can range from 2D optical images to 1D radar range profiles. In each case, the goal is to use the information (primarily invariant geometric information) contained in several images to reconstruct the 3D data.

Classical stereo vision is an obvious example of this and one that has been well studied. The images come from multiple cameras which have been carefully calibrated. The individual images are then registered by identifying feature points in the image produced by one camera with those in the other. To facilitate this registration, the multiple cameras are often placed relatively close to each other. Once registered, the data from the multiple images can be combined with the calibration information to achieve 3D reconstruction.

However, in contrast to stereo vision, less work has been done on 3D
reconstruction using multiple images from a single sensor or using images
from multiple uncalibrated sensors, with perhaps disparate viewpoints,
taken at slightly different times. In the case of a single camera or
sensor we will need the object (or sensor) to be moving so that the
successive images provide added information. In the case of multiple
sensors we will need to know that they are looking at the same object, and
we will need to match feature points between images, but we will not
require calibration of the cameras or assume that the images were taken at
the exact same time.

Our work is a blend of geometric, computational, and statistical techniques
that can be applied to the problem of reconstructing the 3D geometry,
specifically the shape, from multiple images of an object. To have
something concrete, we began by dealing with a collection of feature points
that were tracked from image (or range profile) to image (or range profile)
and we sought to reconstruct the 3D point cloud's shape up to certain
transformations---affine transformations in the case of an optical sensor
and rigid motions (translations and rotations) in the radar case.

Shape in this context has a specific meaning, namely we are looking at the
space (manifold or variety) of n-tuples of points modulo the action of an
appropriate transformation group. Each point in the shape space represents
a geometric configuration of n points up to the allowable transformations.
Two configurations will be the same point in shape space
if they differ by one of the transformations.

In the affine (optical) case the shape spaces for n-tuples of points in
both 3D (object space) and 2D (image space) are familiar Grassmann
manifolds. For a fixed object shape, one can determine all the image
shapes that arise as images of that object using the so-called object/image
equations, which in this case are certain well-known incidence relations.
The resulting locus in the image Grassmannian is a well-known subvariety of
the Grassmannian called a Schubert variety.
Each object shape is shown to be linked to an essentially unique
Schubert variety (image locus) in the image Grassmannian.

Each image of our n feature points will lead to a point in the image shape
space that lies on the Schubert variety (image locus) of our original n-
tuple of 3D points. With enough image points we can fit that locus to
determine the Schubert variety and hence the object shape. The problem
thus becomes one of fitting possibly noisy data in a Grassmannian with a
certain type of submanifold.   In a future paper we plan to explain a
computational approach to this problem.

The radar case is similar. Imagine a rigid collection of n points moving
in space or in the plane. A 1D range profile is constructed by projecting
the points orthogonally onto a line determined by a look vector. The
information in that 1D range profile is essentially the distance to each
point, or more importantly, the relative ranges of the points. We take a
sequence of these range profiles as the point cloud moves in space,
tracking the points from profile to profile.

In this radar case, the image shape space is just an ordinary Euclidean
space of dimension n-1 (translation is removed). Each relative range
profile provides a point in this space. Moreover all the range profiles of
a single 3D (or 2D) n-tuple can be shown to lie on an ellipsoid (or
ellipse) in this space. With enough images we can fit the data in the

image shape space with an ellipsoid (or ellipse), and thereby determine the original 3D (or 2D) shape.

The central issue is how to find the object that best fits the data provided by the accumulated set of noisy images. In our formulation this means finding the object whose image locus in the image shape space best fits the image data. The key is what is meant by "best fit." We argue that the natural Riemannian metric on the image shape space is the best measure to use in determining "goodness of fit," and we are attempting to design an optimal fitting procedure based on this idea for this radar case.

This is not however a trivial problem. Even the simplest case comes down to fitting an ellipse in the plane to a collection of noisy data points. This is a classical problem that has a number of proposed computational approaches, each with different advantages. The difficulty of course is that most of our observations are relatively close together on the ellipse making accurate fitting difficult in the presence of noise.
This is interestingly similar to what one encounters in determining planetary or other orbits from limited observations over a small arc.

Our work on 3D reconstruction in the radar case is joint with Dr. Greg Arnold and Dr. Matt Ferrara at ARFL RYAT. By exploiting shape metrics we hope to achieve optimal 3D reconstruction of configurations of point scatterers. We are also working on an enhanced fitting algorithm (to recover shape from radar range information) to account for the sequential nature of the images (multiple radar range profiles). This involves looking at the tangent bundles to the shape spaces.

2. Shape Statistics

Kendall pioneered statistical shape theory for point features in the plane under similarity transformations. Among his results is a description of the distribution of shapes for point features selected from independent spherical normal distributions each with covariance matrix normalized to the 2 by 2 identity matrix and with means at selected points in the plane. One can regard this as an early attempt to introduce the idea of "noisy" shapes. An important question is to determine for a given distribution of object shapes, the corresponding distribution of image shapes under appropriate hypotheses. This was something not addressed by Kendall or others working in this area. Building on Kendall's results in 2D, we have made progress in answering the above question in a particular case involving a small number of point features in the plane under similarity transformations which are projected to 1D. This is a modified radar case where scale is unknown.

3. Global descriptions of shape spaces in the orthographic (radar) case.

In order to carry out our program of developing the global version of the object/image equations and metrics for the orthographic case and understanding the robustness of our 3D reconstruction algorithms, it is necessary to understand how the shape spaces for point features in this case isometrically embed in standard Euclidean space. For small numbers of point features in 1D this is relatively easy, but for greater numbers of points in 1D and any number in 2D or 3D, this becomes a harder problem. It essentially amounts to finding an embedding of real or complex projective space isometrically into a Euclidean space (real or complex) of as low a dimension as possible, and then extending this embedding to a certain cone

over the projective space. We have been able to do this, paving the way
for the full development of our approach to recognition in the radar case

4. Mapping 2D images to 3D models.

This is a new aspect of our work that has occupied our attention over
the past year.   The problem is a simple one. Imagine you have an object
deployed at a remote location and that you have an accurate CAD model of
the object. The object is subject to surface damage. Patches to repair
the damage can only be manufactured at special facilities located away from
the area of deployment. We want to be able to take a photo of the
damaged area and send it back to the fabrication site so the appropriate
patch can be made. This requires accurately mapping the damaged region to
the CAD model. Generally only a few key features of the object will appear
in the 2D image to help with the alignment.

Our approach assumes that we know the focal length of the digital camera
and details about the CCD, namely its size in millimeters and the number of
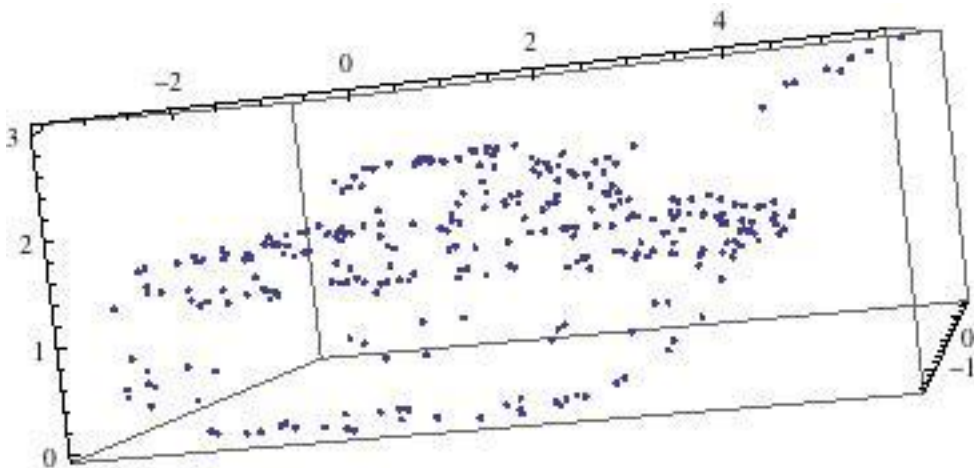pixels in the array.

Using this we can properly align the camera in CAD space and then map
the pixels showing the damaged area to the model using a ray-tracing like
procedure. Our method uses something we developed called "quaternion
optimization".

The algorithm has been implemented in Mathematica and details will be
included in a forthcoming paper for SPIE "Aligning images with CAD models
via quaternion optimization," which will be presented at the annual meeting
in August 2011 in San Diego.

5. Extraction and segmentation of data lying on planar facets within a
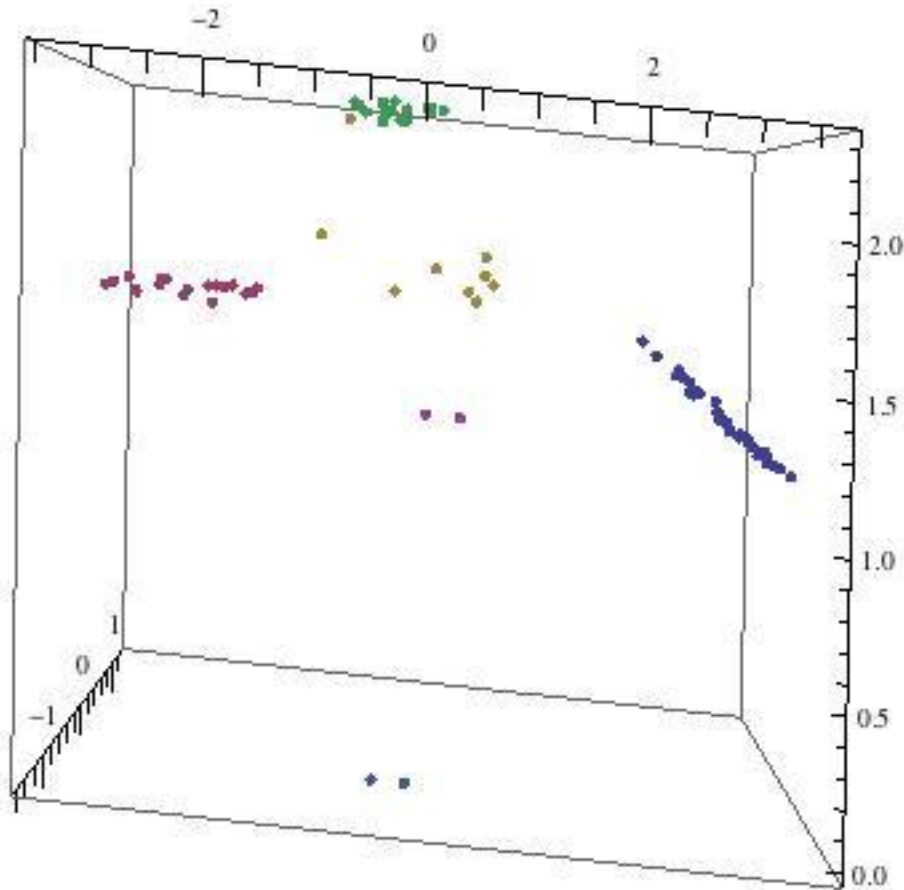   large point cloud.

We have been looking at ways to use GPCA (generalize principal component
analysis) to extract planar features from ladar data. The result is
something we call MGPCA Modified General Principal Component Analysis.
These ideas can be used to align ladar data to CAD models and to speed
various algorithms for matching point clouds to target models. In some
cases it can be used directly to identify targets via an invariant feature
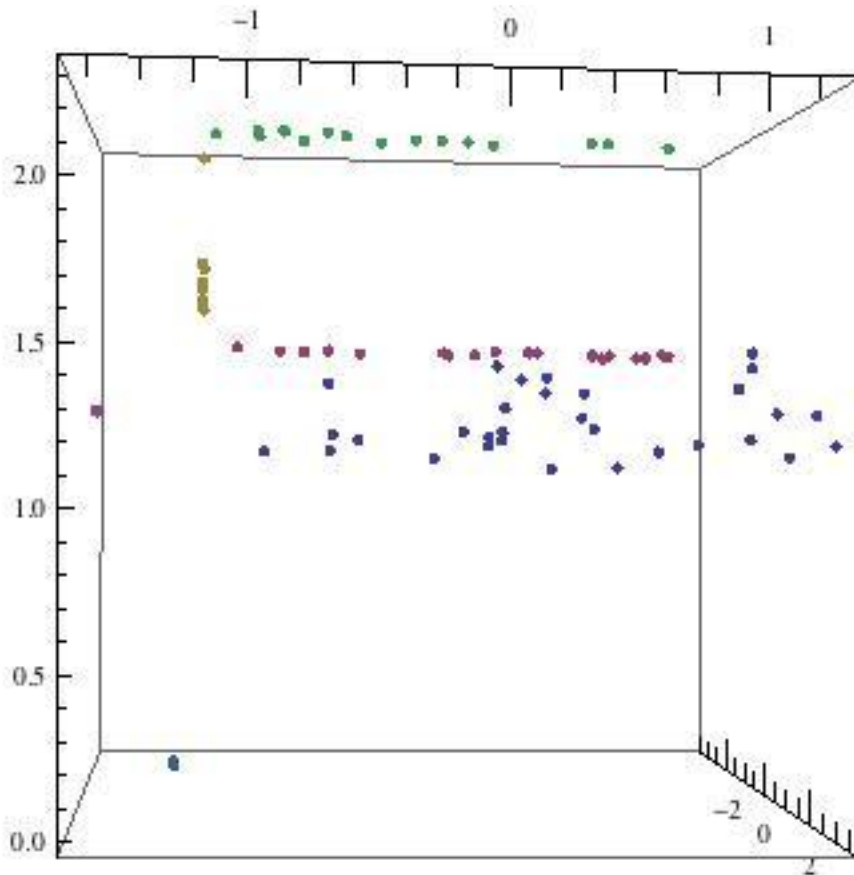based formulation (object/image metric) with planar (instead of point)
features.

Simulated Ladar Point Cloud

Extracted Planar Features

What you see below are points extracted from the point cloud above that lie
on the sloped front of the tank, the rear deck, the top of the turret and
the side of the turret. The extraction and segmentation was done using a
modification of GPCA. One advantage of this method is that planar faces
can be detected even if they are partially obscured. (Note: there is some
distortion in the graphics.)

6. Match a 3D arrangement of point features on an object to an arrangement of 2D point features in an optical image with no a priori knowledge of the camera position, camera parameters, pose, or scale.

We have made use of some recent progress in algebraic geometry to address two fundamental problems related to object recognition for point features under full perspective projection, i.e. a pinhole camera model of image formation. The first involves using recent work of Vakil et. al. to better describe the shape manifolds for arrangements of point features in low dimensional terms. This work uses the representation theory of the symmetric group, outer automorphisms, and recent results in algebraic geometry. The second problem involves determining global equations for the geometric constraints (object/image equations) that must hold between a set of object feature points and any image of those points under a full perspective projection (standard optical camera). These constraints are formulated in an invariant way, so that object pose, image orientation, or the choice of coordinates used to express the feature
point locations on either the object or in the image are irrelevant.
In this recent period of performance we explicitly derived these
equations.

Personnel Supported:

In addition to the principal investigator, the project provided partial support for several graduate research assistants. Our student, Mr. Kevin Abbott graduated in late 2007, and as a result no funds were requested in

this grant for a new student during the first nine months of this project. Entering the second and third year of the project (1 December 2008 to 30 November 2010) funds for a student were available. Mr. Bryan Ko, a senior graduate student in the Mathematics Department worked on the project during part of this time. He received his Ph.D. in January of 2010 under Dr. Stiller's direction. Another graduate student, Corey Irving, was supported during Spring 2011.

Faculty: Dr. Peter F. Stiller, Prof. of Mathematics and Computer Science

Graduate Student: Bryan Ko (received Ph.D. 1/2010)
                 Corey Irving (expected to receive Ph.D. 5/2012)

Publications:

Several publications dealing with this project's results have or will appear in print shortly.

Stiller, P. F., "Aligning images with CAD models via quaternion optimization," Mathematical Methods in Pattern and Image Analysis, SPIE, San Diego, 8/11, to appear.

Stiller, P. F., "Object/Image Relations in Full Perspective and Weak Perspective and 3D Reconstruction," Mathematics of Data/Image Coding, Compression, and Encryption with Applications XII, Proceedings of SPIE, Vol. 7799, San Diego, CA, 10 pages, 8/10.

Stiller, P. F., "Aspects of 3D Reconstruction," Computational Imaging / Mathematical Imaging, Vol. 7246, IS&T/SPIE Electronic Imaging Science and Technology, San Jose, CA, 12 pages, January 2009.

Stiller, P. F., "Object/Image Relations, Shape Spaces, and Metrics – The Generalized Weak Perspective Case," preprint, revised for submission, 30 pages, October 2008.

Mendoza, O. L., Arnold, D. G., and Stiller, P. F., "Further Exploration of the Object-Image Metric with Image Registration in Mind," Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2008, Vol. 6974, SPIE Defense and Security Symposium, Orlando, FL, 12 pages, March 2008.

Interactions/Transitions:

In March 2008, Dr. Stiller met with Dr. Matthew Ferrara and Ms. Olga Mendoza of AFRL/RYAT at the SPIE Defense and Security Symposium in Orlando, Florida. Ms. Mendoza presented our joint paper in the conference on Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications.

In May 2008, Dr. Stiller visited the Air Force Research Laboratory's Target Recognition Branch AFRL/RYAT where plans for collaborative work with Dr. D. Gregory Arnold, Dr. Matthew Ferrara and Ms. Olga Mendoza were made and several of the topics in the proposal were discussed.

In August 2008, Dr. Stiller again visited Drs. Arnold and Ferrara at the Air Force Research Laboratory's Target Recognition Branch AFRL/RYAT. This visit resulted in a joint paper, "Aspects of 3D Reconstruction." During

this visit Dr. Stiller also attended the [ATR Center Summer Workshop,](#) in Beavercreek, Ohio, August 11-14, 2008, which was sponsored by the ATR Center at Wright Paterson AFB.

In January 2009 Dr. Stiller attended the IS&T/SPIE Electronic Imaging Science and Technology Conference in San Jose, California. In the session on Computational Imaging / Mathematical Imaging, he presented a paper entitled "Aspects of 3D Reconstruction." One interesting research contact to come out of this meeting was with a civil engineering group interested in determining if the optimal mix of rock shapes was going into a concrete mixture. This involves a 3D shape reconstruction problem from scanned 2D data. After the SPIE meeting in San Jose, Dr. Stiller attended the Workshop on Classical Algebraic Geometry at the Mathematical Sciences Research Institute in Berkeley, California.

In May 2009, Dr. Stiller again visited the Air Force Research Laboratory's Target Recognition Branch to conduct joint research with Dr. Matt Ferrara on the radar 3D reconstruction problem.

Over the course of the Spring 2009, Dr. Stiller attended several workshops at the Mathematical Sciences Research Insitute in Berkeley, California. These were focused on the lastest developments in algebraic geometry.

Dr. Stiller was an invited participant in the ATR Center Summer Workshop: Performance Modeling and Prediction, held at Wright State University, Beavercreek, OH, August 11-12, 2009.

Following that Dr. Stiller attended the NSF/DOD Planning Meeting for an Industry / University Collaborative Research Center on Surveillance Theory organized by the Air Force Research Laboratory, Wright-Patterson Air Force Base, Dayton, OH, August 13-14, 2009.

In May 2010, Dr. Stiller again visited the Air Force Research Laboratory's Target Recognition Branch to conduct joint research with Dr. Matt Ferrara and Dr. Jared Culbertson AFRL/RYAT on various target identification problems.

Upcoming in August 2011, Dr. Stiller will be chairing the session on "Pattern Recognition Theory" as part of the conferences on Image and Signal Processing at the annual SPIE Optics and Photonics meeting in San Diego, Ca from 8/20/11 to 8/25/11. He will also be speaking in the session on Imaging Theory I – Modeling on "Aligning Images with CAD Models via Quaternion Optimization."

New Discoveries, Inventions, or Patent Disclosures:

Beyond the research results discussed above, there are no new discoveries, inventions, or patent disclosures.

Person completing this report:

Dr. Peter F. Stiller
Professor of Mathematics and Computer Science
Texas A&M University
Phone: (409) 862-2905
Fax: (409) 845-5827
Date: 15 August 2011